

UNIVERSITY OF
ILLINOIS LIBRARY
AT URBANA-CHAMPAIGN
BOOKSTACKS

UNIVERSITY LIBRARY

UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

The person charging this material is responsible for its renewal or return to the library on or before the due date. The minimum fee for a lost item is **\$125.00, \$300.00** for bound journals.

Theft, mutilation, and underlining of books are reasons for disciplinary action and may result in dismissal from the University. *Please note: self-stick notes may result in torn pages and lift some inks.*

Renew via the Telephone Center at 217-333-8400, 846-262-1510 (toll-free) or circlib@uiuc.edu.

Renew online by choosing the **My Account** option at: <http://www.library.uiuc.edu/catalog/>

DEC 30 2008

Digitized by the Internet Archive
in 2011 with funding from
University of Illinois Urbana-Champaign

Hierarchical Spline Models for Conditional
Quantiles and the Demand for Electricity

The Library of the
AUG 5 1990
University of Illinois
of Urbana-Champaign

Wallace Hendricks
Roger Koenker



BEBR

FACULTY WORKING PAPER NO. 90-1661

College of Commerce and Business Administration

University of Illinois at Urbana-Champaign

June 1990

Hierarchical Spline Models for Conditional Quantiles
and the Demand for Electricity

Wallace Hendricks

Roger Koenker

Department of Economics
University of Illinois
Champaign, IL 61820

This research was supported by the Illinois Department of Energy and National Resources and NSF Grant SES-8707169.

ABSTRACT

Methods for estimating nonparametric models for conditional quantiles are suggested based on the regression quantile methods of Koenker and Bassett (1978). Spline parametrizations of the conditional quantile functions are used. The methods are illustrated by estimating hierarchical models for household electricity demand using data from the Chicago Metropolitan Area.

KEY WORDS: Regression quantiles, splines, hierarchical models, non-parametric regression

Hierarchical Spline Models for Conditional Quantiles and the Demand for Electricity

Wallace Hendricks and Roger Koenker

"What the regression curve does is give a grand summary for the averages of the distributions corresponding to the set of x 's. We could go further and compute several different regression curves corresponding to the various percentage points of the distributions and thus get a more complete picture of the set. Ordinarily this is not done, and so regression often gives a rather incomplete picture. Just as the mean gives an incomplete picture of a single distribution, so the regression curve gives a correspondingly incomplete picture for a set of distributions."

Mosteller and Tukey (1977, p.266)

1. Introduction

The classical theory of the linear model as Mosteller and Tukey suggest is essentially a theory for models of conditional expectations. However, in many applications it is fruitful to go beyond these models. The extensive recent literature on estimating models of heteroscedasticity offers one approach. Another, potentially more flexible, approach is to consider models for various conditional quantiles. In strictly linear models a simple approach to estimating models of conditional quantiles is suggested in Koenker and Bassett (1978). In the spirit of non-parametric regression, the work of Stone (1977) and Truong (1989) in the nearest-neighbors tradition, Samanta (1989) and Janssen and Antoch (1989) using kernel based ideas, and White (1990) using neural networks offer other methods of estimating conditional quantile response surfaces.

In this paper we would like to illustrate a new approach to non-parametric estimation of conditional quantile functions employing a linear spline parameterization of the quantile response and the methods of Koenker and Bassett (1978). While splines have proven to be a very flexible tool for nonparametric regression, their application has been generally restricted

to estimating models for conditional central tendencies. See Hastie and Tibshirani (1986) on applications to generalized linear models, Buja, Hastie and Tibshirani and discussion (1989) and Lenth (1977) or Cox (1983) on M -estimation of spline models.

Our application involves modeling daily electricity demand by households in the Chicago metropolitan area. Because of the capital intensive technology required to generate electric power, it is important to understand the determinants of the load cycle, and particularly important to understand these determinants when the load is unusually large, i.e., at high quantiles of the demand distribution. We study this by first fitting parametric models of the daily/weekly demand cycle for electricity for several hundred households, and then (hierarchically) estimating models for the demand parameters in terms of household characteristic like appliance ownership, family size, etc.

We believe there are many other contexts in which similar models may prove useful. In studying pollution data, for example, models for mean concentration levels may be less relevant from a public health standpoint than comparable models for upper quantiles representing more extreme concentration levels. In analysis of standardized test data, trends in mean performance may be usefully supplemented with similar models for other quantiles of the performance distribution. In the econometric literature on the estimation of production technologies there has been considerable interest in estimating so-called "frontier production models" which correspond closely to models for extreme quantiles of a stochastic production surface.

Our paper is organized as follows. Section 2 describes a general approach to estimating non-parametric quantile response functions, Section 3 some associated inference apparatus. Section 4 describes our application including the hierarchical aspect of the model employed. Section 5 discusses data sources, while Section 6 presents the results.

2. Nonparametric Estimation of Conditional Quantile Functions

Estimation of linear models for conditional quantile functions has been considered in a series of papers by Koenker and Bassett (1978, 1982, 1986). The basic idea is exceedingly simple: any θ th quantile of a scalar random variable Y may be viewed as a solution to the problem,

$$\min_{t \in \mathbb{R}} E \rho_{\theta}(Y - t)$$

where $\rho_{\theta}(\cdot)$ is the "check" function $\rho_{\theta}(u) = \theta |u|^{+} + (1 - \theta) |u|^{-}$, illustrated in Figure 2.1.

Figure 2.1




Figure 2.1. Check Function $\rho_{\theta}(u)$ for $\theta = .15$. This function may be used to characterize the θ^{th} quantile as the solution to an optimization problem.

To see this, observe that

$$\frac{d}{dt} [\theta \int_t^\infty (y - t) dF + (1 - \theta) \int_{-\infty}^t (t - y) dF] = F(t) - \theta$$

so at a minimum $F(t) = \theta$. Thus, the θ^{th} sample quantiles from a sample $\{y_1, \dots, y_n\}$ solve

$$\min_{t \in \mathbb{R}} \sum_{i=1}^n \rho_\theta(y_i - t).$$

Similarly in "regression settings" where we might hypothesize a linear relationship between the conditional quantiles of Y and a vector of covariates $x \in \mathbb{R}^p$, i.e.

$$F_Y^{-1}(\theta \mid x) = x' \beta, \quad (2.1)$$

we may define the θ^{th} regression quantiles of the sample $\{(y_i, x_i)\}_{i=1}^n$ as solution to

$$\min_{b \in \mathbb{R}^p} \sum_{i=1}^n \rho_\theta(y_i - x_i' b). \quad (2.2)$$

Some finite sample properties of regression quantiles are discussed in Koenker and Bassett (1978) and their asymptotic behavior is further developed in Ruppert and Carroll (1980), Jurečková (1984), Koenker and Bassett (1986), Koenker and Portnoy (1988), Portnoy and Koenker (1989) and elsewhere.

Most of the existing theory deals with the simple case of linear models with *iid* additive disturbances,

$$y_i = x_i \beta + u_i.$$

In this (leading) case, the regression quantiles have an asymptotic theory that parallels very closely the theory of the ordinary sample quantiles in the one-sample problem. In non-*iid* settings the situation is somewhat more complicated. Consistency is essentially ensured by (2.1) and a mild design condition on the growth of $\sum x_i x_i'$ and positivity of the conditional density, see Koenker and Bassett (1986). Portnoy (1990) has recently proved a linear representation theorem for regression quantiles that implies corresponding asymptotic normality results under quite general heterogeneity and dependence conditions. White (1990) has also studied

estimation of conditional quantile functions under quite general conditions using a neural network approach.

In this paper we are particularly concerned with situations involving quite arbitrary heterogeneity in the conditional distribution of the response variable over the design space. To fix ideas consider a single design variable x , and suppose (only) that the conditional quantiles of a response variable Y are smooth in x , so for suitable basis functions $\{\phi_i(x)\}_{i=1}^{\infty}$ we may approximate *cqf's* of Y by

$$F_Y^{-1}(\theta \mid x) = \sum_{i=1}^p \beta_i(\theta) \phi_i(x). \quad (2.3)$$

Obviously, the choice of $\{\phi_i(x)\}$ is application dependent. But cubic splines offer many natural advantages for a broad class of problems. Models like (2.3) with a judicious choice of a few spline $\phi_i(\cdot)$'s are sometimes called "parametric splines" or "regression splines" to distinguish them from the "smoothing spline" models championed by Wahba and others. Stone (1985) makes a strong case for such parametric splines. See also Ramsey (1988) and the discussion thereof regarding the relative merits of parametric and smoothing splines.

Ramsey (1988) also makes a strong case for the convenience of spline models for situations in which *a priori* monotonicity constraints need to be imposed as well as smoothness. A simple approach to isotonic splines is available by simply treating the integrated B-splines as basis functions and restricting the associated coefficients to be positive. Convexity can also be imposed through another integration. We do not explore this approach here, but we might note that solving (2.2) subject to nonnegativity constraints on the parameter vector requires only a slight computational modification since the unconstrained problem is a linear program. The discussion of Ramsey also contains a number of important cautionary remarks regarding *a priori* monotonicity restrictions.

In our application it is important to have *periodic* splines, since we will be estimating hourly models of daily and weekly load cycles for electricity demand. This too is straight-

forward. We simply require that two ends of the spline join smoothly in the same fashion as at the other knot locations. See e.g., DeBoor (1978). In the case of cubic splines, used here, this means the function and its first two derivatives are continuous at the endpoints.

The linear parametrization of the regression function afforded by the spline formulation is a significant advantage since it makes estimation a straight forward exercise in linear programming. See Koenker and D'Orey (1989) and Osborne (1990) for detailed descriptions of algorithms for the linear regression quantile problem.

Alternative approaches based on nearest neighbors: Stone (1977) and Truong (1989) or kernels: Samanta (1989) and Antoch and Janssen (1989), are computationally more complex since they require separate computations at each design point where an estimated quantile is required. When the dimension of the design space is moderate, say greater than 2, kernel and NN methods appear problematic. However, additive spline models like those discussed in Buja, Hastie, and Tibshirani (1989) seem to offer a tractable, yet flexible approach to multivariate, non-parametric conditional quantile estimation.

Strong consistency of regression quantiles under mild regularity conditions on the design and the linear specification (2.3) is established in Bassett and Koenker (1986). In Appendix A we sketch the argument for the asymptotic normality of $\hat{\beta}_n(\theta)$ under mild conditions. In the next section we take up the problem of hypothesis testing in the context of regression quantile estimation.

3. Wald Tests for Regression Quantile Models

A critical aspect of any estimation scheme for conditional quantile models is the capacity to do formal testing and model selection. Here we will briefly describe a Wald approach. See Koenker and Bassett (1982) and Koenker (1986) for further details and Gutenbrunner, Jurečková and Koenker (1990) for an alternative approach.

When the linear model errors are *iid* there is a well developed asymptotic theory leading to the construction of tests. In this case if the error distribution, F , has strictly positive density at the θ^{th} quantile, i.e. $f(F^{-1}(\theta)) > 0$, and the design satisfies $\lim n^{-1} X'X \rightarrow D$ with $x_{i1} = 1$ for all i , then

$$\sqrt{n}(\hat{\beta}_\theta - \beta_\theta) \rightsquigarrow N(0, \omega_\theta^2 D^{-1})$$

where $\beta_\theta = \beta + (F^{-1}(\theta), 0, \dots, 0)'$ and $\omega_\theta^2 = \theta(1 - \theta)/f^2(F^{-1}(\theta))$. Thus if we are interested in a test of

$$H_0: R\beta = r$$

it is natural to base the test on a statistic of the form

$$\xi = \hat{\omega}^{-2}(R\hat{\beta}_\theta - r)'(R(X'X)^{-1}R')^{-1}(R\hat{\beta}_\theta - r).$$

where $\hat{\omega}$ denotes some consistent estimator of ω . Note that the precision with which the parameters of the θ^{th} quantile function are estimated is inherently controlled by the magnitude of the density at this quantile. Thus quantiles in the tail of the distribution where the density is low are inherently more difficult to estimate and therefore corresponding tests have reduced power relative to quantile whose density is higher.

It remains to consider the estimation of the nuisance parameter ω . The essential feature of this problem is to estimate the so-called sparsity function

$$s(\theta) = 1/f(F^{-1}(\theta)).$$

Siddiqui (1960) suggested, in the one-sample $\{X_1, \dots, X_n\}$ model,

$$\hat{s}_n = \frac{n}{2d_n} [X_{(\lfloor n\theta \rfloor + d + 1)} - X_{(\lfloor n\theta \rfloor - d + 1)}] \quad (3.1)$$

Where $X_{(i)}$ denotes the i^{th} order statistic from $\{X_1, \dots, X_n\}$. This is a kind of histogram method and various other approaches are possible. See e.g., Koenker and Bassett (1982) and Welsh (1987). Hall and Sheather (1988) have intensively investigated the Siddiqui approach and we restrict attention to this approach here.

Standard density estimation asymptotic considerations, see Sheather and Maritz (1983), suggest that the "bandwidth" d_n in (3.1) should be $d_n = d_0 n^{4/5}$ where $d_0 = (9s^2(q)/s''(q)^2)^{1/5}$. In Figure 3.1 we illustrate optimal D_0 's for 3 quite different distributional shapes. We have used the "Normal" d_0 's in the tests reported below.

Figure 3.1




Figure 3.1. Optimal Siddiqui Constants for 3 Densities. Optimal constants for the Siddiqui bandwidth parameter are plotted against the quantile of interest. Note that the choice of constant is invariant to location and scale of the underlying density.

In the *iid* case this approach may be implemented simply by replacing the order statistics in (3.1) by the corresponding order-statistics of the residuals from the θ^{th} regression quantile fit. However, in non-*iid* circumstances this leads to inconsistent estimates of the covariance matrix of $\hat{\beta}_\theta$. See Appendix A for details showing that in the independent not identically distributed case

$$\sqrt{n} (\hat{\beta}_\theta - \beta_\theta) \rightsquigarrow N(0, B_n^{-1} A_n B_n^{-1})$$

where $A_n = \theta(1 - \theta)X'X/n$ as above, and

$$B_n = n^{-1} \sum_{i=1}^n f_i(F_i^{-1}(\theta))x_i x_i'$$

where f_i, F_i denote the marginal density and *cdf* of the i^{th} error observation respectively. Estimating B_n may be accomplished in several ways. Our proposal, following the sparsity estimation literature closely, is as follows. Take d_n as defined above and compute $\hat{\beta}_\theta$ for $\theta^\pm = ([n\theta] \pm d_n + 1)/n$. Then at each design point compute

$$\hat{f}_i = 2d_n/(nx_i(\hat{\beta}_{\theta^+} - \hat{\beta}_{\theta^-}))$$

and finally

$$\hat{B}_n = n^{-1} \sum \hat{f}_i x_i x_i'.$$

Under rather mild conditions, one can show that $\hat{\beta}_n \rightarrow \beta$ in probability. Thus the procedure provides a regression quantile analogue to the popular Eicker-White heteroscedasticity-consistent covariance matrix for the least-squares estimator. A potentially troublesome feature of the \hat{f}_i 's is the fact that they *may* be negative. We know (Bassett and Koenker (1982)), Thm 2.1) that $\bar{x}'\hat{\beta}(\theta)$ is monotone in θ , but at $x_i \neq \bar{x}$ this is not assured. How important this is in practice remains to be seen.

4. Hierarchical Models for Conditional Quantiles and Their Application to Household Electricity Demand

Virtually all of the research on statistical models of household electricity demand employs some variant of the hierarchical linear models introduced by Lindley and Smith (1973). See e.g., Hendricks, Koenker and Poirier (1979) and Engel, Granger and Weiss (1986). The common data structure of such research is a long, high frequency time series on a large sample of individual households--in our present case we have an hourly time-series for the four summer months of 1985 for about 400 households. The analysis typically proceeds in two

stages: in the first, a model of the demand cycle is estimated for each household effectively collapsing a time-series of several thousand observations into a few estimated parameters; and in the second stage, a model is specified to explain the cross-sectional variability of these demand cycle parameters based on various demographic and economic characteristics of the households. An excellent exposition of the statistical foundations of such hierarchical models from a Bayesian standpoint may be found in Smith(1975).

4.1. Stage I: A Time-Series Model of the Household Demand Cycle

In this section we will describe in detail the parametric model that we have used to characterize household demand behavior. It is convenient for purposes of interpretation to decompose our model into two components: one which is strictly periodic and therefore insensitive to the weather, and another so-called weather-sensitive component. We will describe these two components in turn.

Let $y(t)$ denote demand in kw at time t . A strictly periodic model of demand may be written in the form

$$y(t) = \sum \alpha_i x_i(t) + u(t) \quad (4.1)$$

where $x_i(t)$: $i=1, \dots, p$ denote functions periodic on some time interval like a day, or week. The error process $u(t)$ will be discussed further below. There are obviously many possible competing choices for the functional form of the $x_i(t)$'s. In several studies they are the sines and cosines of classical Fourier series. In our earlier work cubic splines were used. We tend to prefer the spline formulation because the coefficients have a ready interpretation as fitted demand at particular times-of-day, i.e. the knot positions. We do not regard the choice of harmonic functions as crucial; many families would serve adequately and therefore the choice is essentially one of computational and interpretative convenience.

In our application there are potentially two frequencies at which demand exhibits periodic behavior: daily and weekly. To accommodate both frequencies we construct two

periodic splines. The first is a daily spline with 4 knots at midnight, 6 am, noon, and 6 pm, observed hourly; and the second is a weekly spline with knots at midnight Sunday, midnight Monday and midnight Friday observed at daily frequency. Let D denote the former 24×4 matrix of daily spline effects -- the $x_i(t_j)$ $i=1,\dots,4$, $j=1,\dots,24$ in (4.1) above, and W denote the 7×3 matrix of the weekly spline effects. The design matrix for the full weeks non-weather sensitive component is then simply

$$S = W \otimes D$$

which is a 168×12 matrix. Thus there are 12 parameters that characterize the strictly periodic component of the demand model.

The weather sensitive component of demand is potentially more controversial since, at least potentially, many factors might be thought to influence this component. After considerable exploratory investigation on a small subsample of households, we have returned to a specification that we have used in the past that might be roughly characterized as a "ratchet temperature effect" model. The basic idea is quite simple. Temperature is presumed to influence demand through its current level, its maximum level in the last day, and its maximum level in the last two days. Thus a sequence of hot days may generate a different demand response than an isolated hot day. This seems to accord with casual empiricism, and worked well in our preliminary exploration. Since the households in the sample are scattered around the Chicago area each household is matched with the nearest of 6 weather stations on which we have a full time-series of hourly temperatures for the summer of 1985. The addition of these three temperature variables yields a model of the demand cycle with 15 parameters. Obviously, one might consider more complex models in which the weather sensitive component and the non-weather sensitive (strictly periodic component) were not simply additive. However, given the noise level in the first stage of the model it seems unlikely that such extensions would yield informative results. Similarly, experimentation with other aspects of weather like humidity failed to improve upon the simple temperature specification suggested

above.

4.2. Stage II: A Cross-Section Model of the Demand Cycle

The second stage of the analysis involves explaining variation in the first stage parameters across households. Thus the model takes the form of the hierarchical models of Lindley and Smith (1973) and primary interest focuses on the meta-parameters of the second-state model that describe in detail the effect of household characteristics (appliance holdings, etc.) on the shape and level of the household demand cycle. It is at this stage that the analysis offers a statistical decomposition of the average load cycle into distinct demand contributions by end-use.

Formally, we may express the first stage parameters, represented by the p -vector α as linear functions of household characteristics in the multivariate linear model

$$\alpha_i = z_i B + v_i \quad (4.2)$$

where α_i is a p -vector of parameters representing the level and shape of the demand cycle incorporating perhaps its responsivity to weather variables, z_i is a k -vector of household characteristics, B is a $k \times p$ matrix of coefficients and v_i is a random p -vector. The meta-parameters B are of fundamental importance to any end-use analysis of load-shapes since they provide the connection between household composition/appliance stocks and the level and shape of the household load cycle.

One might wonder, encountering these models for the first time, why not simply substitute (4.2) into (4.1) and proceed directly to estimate the parameters B . However the dimensionality of the resulting estimation problem is awe inspiring, and, making the eminently plausible assumption of independence across households in demand behavior, the problem is most conveniently dealt with in two explicit stages.

Presuming that the first stage model (4.1) can be estimated efficiently, we are then faced with the problem of estimating the Stage II model (4.2). Here the problem begins with the

obvious fact that we do not observe α_i directly, instead we have $\hat{\alpha}_i$, the *estimated* parameters describing the load cycle of the i^{th} household. Thus (4.2) becomes

$$\begin{aligned}\hat{\alpha}_i &= z_i B + (\hat{\alpha}_i - \alpha_i) + v_i \\ &= z_i B + w_i.\end{aligned}\tag{4.3}$$

The new error is composite: partially attributable to v_i the original Stage II error and partially due to the Stage I error in estimating the vector α . Again, the resulting complexity of the error specification creates certain complications of estimation. The error component $(\hat{\alpha}_i - \alpha_i)$ can be ignored completely and estimation of the second stage then proceeds as if $\hat{\alpha}_i = \alpha_i$. Here we experimented with weighting the cross-sectional observations by the reciprocal of the standard error of the Stage I fits. However, this seemed to have the effect of downweighting households with large demand and thus produced less reliable results than the unweighted estimates. The success of such simple weighting methods must ultimately depend on the scale of the two error components in (4.3) being proportional in the sample; this seems quite implausible and therefore more sophisticated methods are called for. As in previous work, we have found that the variance of v in (4.3) is substantially greater than the estimation error component and thus little is lost due to the omission of the weights.

The choice of the exogenous variables z_i is obviously highly application dependent. In the present instance we have a somewhat limited inventory of appliance stock variables; section 5 provides a complete description of these data and Section 6 interprets the Stage II results.

4.3. Hierarchical Models for Conditional Quantiles

The two-stage hierarchical framework sketched in the previous section has traditionally been viewed as a means to specify a model for conditional expectations. The Stage I demand model has been estimated by classical least squares methods and consequently may be viewed as an estimate of the conditional mean of the demand cycle under given temperature conditions. However, it seems implausible to assume that the error $u(t)$ in (4.1) is stationary. The

stationarity assumption on $u(t)$ is extremely strong; formally, it requires that the joint distribution function of the $u(t)$ evaluated at several points in time, say t_1, \dots, t_k , remains identical if we shift the time points forward or backward uniformly in time. Thus, *a fortiori*, moments of $u(t)$ are independent of t , and the dispersion, or skewness, of the variability of demand around its central tendency is *assumed* constant. It is clear, however, that this assumption is highly implausible. Like the cyclical behavior of the mean, the stochastic component of demand is undoubtedly also highly cyclical. The dispersion and skewness of demand is likely to be higher when mean demand is high, and lower when demand is lower. This is, of course, most clear in extremely low demand periods in which constant variability in demand would threaten the obvious physical necessity that demand be non-negative.

A general approach to the specification and estimation of models of the load cycle is afforded by the regression quantile methods. Rather than specifying a periodic model for the mean demand cycle and regarding noise as a necessarily stationary process added to this central tendency, one may specify several models for distinct quantiles of the load cycle and analyze their behavior as potentially distinct phenomena. Thus, for example, it seems reasonable to view households as having a "baseload" demand which is quite stable corresponding to, say, the 10th percentile of the demand cycle. (Stable in this context effectively means that there is little difference between the lower quantiles, or to put it yet another way, that the load cycle distribution has a short left tail.) A median demand cycle might be specified in much the same way as one specified the mean-cycle model of (4.1). Finally, a model for some large quantile, say the 90th percentile, would reflect the behavior of the upper extreme of the stochastic load cycle.

In Figures 4.1-2 we present estimated conditional quantiles of the the demand cycles for the first and second household in the sample. The fitted quantile functions in these figures are based on the weather conditions prevailing during the first week of June as measured by their closest weather station. If the stochastic component of demand were additive and stationary

we would expect to see in these figures conditional quantiles which were simply vertical displacements of each other. Obviously, this is not the case and it is quite important to distinguish the shape and amplitude of the various quantile estimates. This point seems particularly important in hierarchical models since we naturally associate certain Stage II effects (appliances, for example) with contributions to the baseload, and others to peak demand. In the mean-plus-stationary-noise model such distinctions are meaningless. We will have more explicit remarks to offer along these lines in Section 6 when we discuss the results of the Stage II estimation in detail.

Figure 4.1





Figure 4.2



Figures 4.1-2. Estimated conditional Quantiles of the demand Distribution for Two Representative Households. The plotted curves indicate, in ascending order the .25, .50, .75 and .95 fitted quantiles of the weekly demand distribution measured in kilowatts. On the time axis 0 represents 0:00 a.m. on Sunday; 7 represents 24:00 on Saturday.

5. Description of the Data

In April and June of 1985, the Strategic Analysis Department of Commonwealth Edison conducted a mail survey of 1000 specially-metered Residential Load Study (RLS) customers. The primary goal of the survey was to compare the appliance ownership in the RLS sample to the estimated holdings of the general residential population of Commonwealth customers. The one-page questionnaire was returned by 689 of the original 1000 households.

The RLS consists of magnetic tape demand recorders for 1000 customers located throughout the Commonwealth service area. We obtained KWH data for each month from June through September 1985. In addition these data contain a site indicator matching the data to the six weather stations and an account number matching the data to the survey data. Due to meter malfunctions and outages there are typically 600 to 800 customers monitored each month. We excluded customers who did not have at least 14 days of usable data each

month.

The match of the KWH data and the RLS data created a data set with 371 customers. From this set, we eliminated 30 customers due to missing data on one or more of the nine characteristics that were used in the stage II analysis. Thus the final data set for stage II contained 341 observations.

The weather data consist of hourly measurements on temperature and relative humidity at six sites in Chicago. There were three weeks of missing humidity data for one of the six sites. We reconstructed these data using regression analysis and the humidity data and temperature data from all the sites. However, the temperature and relative humidity series run together extremely well for this time period (a correlation in excess of .9). We therefore decided to specify the stage I model as a function of temperature alone.

5.1. Stage I Results

Stage I estimates are based on fitting the entire summer's data (2880 hourly observations for households with complete data) and plotting the fitted demand cycles for each of the four quantiles: 25th, 50th, 75th and 95th. In addition to these estimates of the conditional quantiles of the demand cycle, we also estimated for each household the conditional mean cycle by traditional least-squares methods. To evaluate the performance of these models it is useful to examine some relevant test statistics. In Table 5.1 we compute median F statistics over the sample of 371 households for four hypotheses. The first hypothesis under test is that there is no daily effect, or more precisely that there is no significant difference in the coefficients of the spline at each of the daily knots. This hypothesis has 9 degrees of freedom and critical value of 1.8 at the 5 percent level. Thus the median test statistic rejects the null except for the case of the 95th percentile. The weekly effect has eight degrees of freedom--the four daily spline coefficients restricted to be the same at each of the three weekly spline knots--and the test statistics are somewhat weaker, but similar to the daily effect. The weather effect is considerably stronger. Here we are testing whether the three coefficients on the temperature

variables are all zero, and the median test statistics are all above the critical value of about 2.6, except again in the case of the 95th percentile. Finally, the last row of the table labeled "spline effect" corresponds to a test of the joint hypothesis that there is no daily or weekly effect, that is that the cycle is entirely weather driven without any strictly periodic component. Again this hypothesis is strongly rejected by the data.

In Table 5.2 we report the proportion of households for which the F statistics discussed above reject the relevant null hypothesis. It is clear from this table that even in the last column the proportion of households yielding significant evidence against the nulls is quite substantial. Details of the procedures used to compute the test statistics for the quantiles are provided in Koenker(1986) and Koenker and Portnoy(1988). We might note here that the variance of the quantile estimates is inversely proportional to the square of the error density at the relevant quantile, and therefore precision is necessarily less in the tails of the distribution than the center for densities of conventional (unimodal) shape.

It is interesting to note that despite the low precision of the estimation of the 95th conditional quantile function, our ability to "explain" the variability of these estimates across

Table 5.1

Medians of the F statistics from Stage I Estimation					
	mean	25%	50%	75%	95%
daily effect	8.49	192.27	60.93	14.10	1.85
weekly effect	6.02	135.26	37.43	9.33	1.14
weather effect	15.77	263.84	83.95	19.15	1.77
spline effect	10.68	262.37	74.99	18.42	2.13

Table 5.2

Proportions of Rejections of F-tests					
	mean	25%	50%	75%	95%
daily effect	0.82	0.97	0.97	0.88	0.47
weekly effect	0.74	0.97	0.97	0.80	0.38
weather effect	0.77	0.98	0.92	0.77	0.45
spline effect	0.86	0.98	0.98	0.90	0.55

households in the next stage of the analysis is really quite impressive. We now turn to this task.

6. Stage II Results

6.1. Specification

Our first stage representation of the load cycle consists of a daily cycle, a weekly cycle, a contemporaneous weather effect and a lagged weather effect. In principle it might be possible to model each of these effects differently and therefore incorporate different explanatory variables for different coefficients in the second stage of the analysis. The major possibility would be to use weather related appliances (e.g. central air conditioning) only for the weather effect coefficients and to exclude non-weather related appliances (e.g. TV sets) in the analysis of the weather effect.

We rejected this strategy for two reasons. First, it is not entirely clear which characteristics only relate to the weather sensitive load and which characteristics only relate to the non-weather sensitive load. For example, it is possible that customers use their electric stove differently on hot days. Second, the daily and weekly cycles are not purely weather insensitive. Since there is no intercept in the first stage, the spline coefficients include both the constant effects of the weather sensitive and non-weather sensitive loads (although it must be remembered that these constant effects are measured at zero degrees). Thus, the impact of weather related characteristics on the constant weather effect can only be captured by their inclusion in the regressions for both the daily-weekly cycle spline coefficients and the temperature coefficients.

Our Stage II model is therefore the same for all the Stage I coefficients and for all quantiles. Expressing the first stage parameters as the p-vector α_i , the second stage can be written as

$$\alpha_i = z_i B + v_i$$

where z_i is a k -vector of household characteristics, B is a $k \times p$ matrix of coefficients and v_i is a random p -vector. We could not use all the information obtained in the RLS demographic survey due to missing observations on some characteristics. In addition some characteristics were almost constant in the sample population. We therefore selected ten characteristics from this survey for inclusion in z . Their definitions are as follows:

FAMILY - the sum of adults and children in the household;
ROOMNUM - number of room air conditioners in the household;
REFRIG - number of refrigerators in the household;
APT - dummy variable for multi-family dwelling;
DISH - dummy variable for presence of dishwasher;
DRYER - dummy variable for presence of electric dryer;
STOVE - dummy variable for presence of electric stove;
WATER - dummy variable for presence of electric water heater;
CENTRAL - dummy variable for presence of central air conditioning;
TV - the number of TV's in the household.

6.2. Estimation Procedure

After considerable experimentation with various weighting schemes the Stage II estimation was carried out by classical unweighted least-squares methods. There are several reasons why weighting might be relevant. As we noted in Section 4, the error in the Stage II equation (4.3) is composed of an estimation error from Stage I and a structural disturbance. The latter we conclude is substantially larger, and while we have estimates of the former quantity there is little reason to suspect that the two error components are proportional, or even positively correlated. Thus weighting household observations by the standard error of the stage I fits yields results which appear considerably worse than those reported here based on unweighted estimation of Stage II.

Another possible rationale for weighting comes from suspected correlation between the errors in the various Stage II "equations". This is the well-known "seemingly unrelated regression" framework; however, here we have common designs in each of the "equations" and consequently there is no efficiency gain from the reweighting. Since even in the second stage

estimating the transformed model as a system is rather unwieldy we chose to estimate it an equation at a time.

6.3. Results

There are 15 Stage I parameters to be estimated for each quantile and for the mean consumption. Since the resulting 75 parameters and their associated errors tend to blur even to the most interested eye, we have relegated them to an appendix available from the first author. We will provide a description of the pattern of the results instead.

There are several strong patterns in the results. First, with the exception of the results for lagged maximum temperature (hereafter LMT), the Stage II variables are able to explain a larger portion of the variation in the extreme (95% quantile) usage than for other quantiles of the distribution. The 95% quantile fit is also slightly better than for mean consumption (again excepting LMT). In addition, the explained variation increases consistently from the lowest to the highest quantiles of the distribution (excepting LMT). Most of the variation in minimum electricity usage by customers can not be explained by our measured characteristics. The R^2 's for the 25% quantile range from 6 to 12%. Thus, minimum electricity usage (baseload?) is probably primarily influenced by behavior rather than the characteristics measured in our survey. As the quantiles increase, our ability to explain their variation also increases. Thus, large increases in usage are much more explainable by our characteristics. This suggests that customers with large numbers of electrical appliances will have much larger variations in their typical usage than that experienced by small users.

Second, the fits for the daily-weekly spline coefficients are essentially all the same. The R^2 's for the mean range from 31 to 35%; those for the 25% quantile range from 10 to 12%; those for the 50% quantile range from 16 to 19%; those for the 75% quantile range from 24 to 30% and those for the 95% range from 32 to 40%. There is almost no detectable weekly cycle in the fits. The daily cycle seems to fit slightly better in the late afternoon-early evening, but this difference is marginal at best.

Finally, the fits for the daily-weekly spline parameters are better than the fit for the contemporaneous temperature parameter and significantly better than the fits for the lagged temperature parameters. In fact, the R^2 values for the two day lagged temperature are not significantly different than zero. Differences across customers in the lagged temperature coefficients are not well explained by the Stage II characteristics.

6.3.1. Patterns for Individual End Use Characteristics

While the coefficients estimated in Stage II provide information on the end-use load curves for our characteristics, they are often difficult to interpret both because the impact of a characteristic for a given hour is a combination of its impact on several of the second stage coefficients and because these impacts sometimes have opposite signs. For example, the effect of central air conditioning is negative for the daily-weekly spline coefficients and positive for the weather coefficients. What then is its total impact at a particular time of day?

To answer this question, we decided to employ a relatively standard technique in the analysis of electricity load curves -- namely dividing the load into "weather sensitive" (WS) and "non-weather sensitive" (NWS, or what we also call baseload) components. The question then became how to define these loads within the context of our estimation procedure.

Our first step was to define a typical week of weather. Since we did not have access to historical data, we took the approach of defining a typical week from our data set. The 1985 summer season was quite cool, so our typical week or hot week are probably cooler than normal in Chicago. The temperature at each hour during our four month period was averaged over the six temperature sites. This temperature series was then fit to a model that included only the 12 daily-weekly spline variables that we used in our Stage I analysis. The model was fit for 10%, 50% and 90% quantiles. The predicted values from these three quantiles were then used to define the following:

Baseload (Cool Week)- This was defined as the predicted weather week of temperatures which are warmer than only 10% of the temperatures during our four months.

Typical Week- This was defined as the predicted weather from the 50% quantile estimates. This week represents a week of temperatures which are warmer than 50% of the temperatures during our four months.

Hot Week- This was defined as the predicted weather from the 90% quantile estimates. This week represents a week of temperatures which are warmer than 90% of the temperatures during our four months.

Non-weather Sensitive Load- This was defined as the load which would occur during a "cool week", e.g. the load estimated with temperatures from the 10% quantile estimates.

Weather Sensitive Load- This was defined as the difference between the load which would occur during the week being analyzed (e.g. a "hot week" or a "typical week") and the load which would occur during a "cool week". Thus, by definition, the weather sensitive load is zero during a cool week.

The maximum temperature variables were defined by finding the appropriate maximums of the given series (10, 50 or 90% quantile estimates). The combination of the given week temperatures and maximum temperatures thus allow us to predict the total and baseline or weather and non-weather sensitive loads for different quantiles as well as to break these loads down into the incremental impacts of various characteristics.

Tables 6.1 through 6.3 give the estimated hourly load by characteristic for a typical "cool" day, a typical summer day and a typical "hot" day. Again, these days are typical for the summer of 1985, but are probably cooler on average than corresponding summer days for other years. These tables provide the estimated total effect of each characteristic by hour for the fit to mean usage. The corresponding values for the 95% quantile are given in Tables 6.4 through 6.6.

Table 6.1

Appliance Load for Typical Cool Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.13	-0.08	0.24	-0.09	0.17	0.29	0.20	0.26	-0.19	0.0
2	0.10	-0.09	0.23	-0.11	0.17	0.25	0.21	0.28	-0.24	-0.02
3	0.08	-0.10	0.23	-0.13	0.17	0.21	0.22	0.31	-0.29	-0.03
4	0.06	-0.10	0.23	-0.16	0.16	0.17	0.23	0.35	-0.34	-0.04
5	0.05	-0.11	0.22	-0.18	0.15	0.15	0.24	0.40	-0.37	-0.05
6	0.05	-0.11	0.22	-0.19	0.14	0.15	0.23	0.44	-0.39	-0.05
7	0.05	-0.11	0.23	-0.19	0.12	0.18	0.22	0.48	-0.38	-0.04
8	0.07	-0.10	0.23	-0.17	0.10	0.23	0.20	0.51	-0.35	-0.03
9	0.09	-0.09	0.24	-0.15	0.09	0.28	0.17	0.54	-0.32	-0.02
10	0.12	-0.08	0.24	-0.14	0.07	0.34	0.14	0.56	-0.27	0.0
11	0.14	-0.07	0.25	-0.12	0.06	0.40	0.12	0.58	-0.22	0.01
12	0.17	-0.07	0.25	-0.11	0.06	0.44	0.11	0.59	-0.18	0.02
13	0.18	-0.06	0.25	-0.12	0.07	0.46	0.10	0.59	-0.13	0.03
14	0.19	-0.07	0.25	-0.13	0.09	0.47	0.10	0.58	-0.10	0.04
15	0.20	-0.07	0.25	-0.15	0.11	0.47	0.11	0.57	-0.08	0.04
16	0.20	-0.08	0.26	-0.16	0.14	0.46	0.13	0.55	-0.06	0.04
17	0.20	-0.09	0.26	-0.17	0.16	0.44	0.14	0.52	-0.06	0.04
18	0.19	-0.09	0.25	-0.17	0.18	0.43	0.15	0.48	-0.05	0.04
19	0.19	-0.09	0.25	-0.16	0.19	0.42	0.16	0.44	-0.05	0.04
20	0.19	-0.09	0.25	-0.14	0.20	0.41	0.17	0.39	-0.05	0.03
21	0.18	-0.08	0.25	-0.12	0.20	0.39	0.17	0.34	-0.07	0.03
22	0.18	-0.08	0.25	-0.10	0.19	0.38	0.17	0.30	-0.08	0.02
23	0.16	-0.07	0.25	-0.08	0.18	0.36	0.18	0.27	-0.11	0.02
24	0.15	-0.07	0.24	-0.08	0.18	0.33	0.19	0.26	-0.15	0.01
sum	3.33	-2.04	5.83	-3.32	3.37	8.10	4.05	10.61	-4.53	0.12

Table 6.2

Appliance Load for Typical Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.18	-0.02	0.13	-0.06	0.16	0.30	0.09	0.10	0.13	0.06
2	0.16	-0.03	0.12	-0.08	0.16	0.25	0.11	0.12	0.08	0.04
3	0.14	-0.04	0.11	-0.11	0.15	0.21	0.12	0.15	0.04	0.03
4	0.12	-0.04	0.11	-0.13	0.15	0.18	0.12	0.18	0.0	0.02
5	0.11	-0.05	0.10	-0.15	0.14	0.16	0.12	0.22	-0.03	0.02
6	0.10	-0.05	0.10	-0.16	0.13	0.16	0.12	0.26	-0.04	0.01
7	0.11	-0.05	0.10	-0.15	0.11	0.19	0.10	0.30	-0.03	0.02
8	0.13	-0.04	0.11	-0.14	0.09	0.23	0.08	0.34	-0.01	0.03
9	0.15	-0.03	0.12	-0.13	0.07	0.29	0.06	0.37	0.02	0.04
10	0.17	-0.02	0.13	-0.11	0.06	0.35	0.03	0.39	0.07	0.06
11	0.20	-0.01	0.13	-0.09	0.05	0.40	0.01	0.41	0.11	0.07
12	0.22	-0.01	0.14	-0.08	0.05	0.44	0.0	0.42	0.16	0.08
13	0.24	-0.01	0.14	-0.09	0.06	0.47	-0.01	0.42	0.20	0.09
14	0.25	-0.01	0.14	-0.10	0.08	0.48	-0.01	0.41	0.24	0.10
15	0.25	-0.01	0.13	-0.11	0.10	0.48	-0.01	0.39	0.27	0.10
16	0.26	-0.02	0.13	-0.13	0.13	0.46	0.0	0.37	0.29	0.11
17	0.26	-0.02	0.13	-0.14	0.15	0.45	0.01	0.33	0.31	0.11
18	0.26	-0.02	0.13	-0.14	0.17	0.43	0.03	0.30	0.31	0.10
19	0.25	-0.02	0.13	-0.13	0.18	0.42	0.04	0.25	0.31	0.10
20	0.25	-0.02	0.13	-0.11	0.19	0.41	0.05	0.21	0.30	0.10
21	0.24	-0.02	0.13	-0.09	0.19	0.40	0.05	0.17	0.28	0.09
22	0.23	-0.02	0.13	-0.07	0.18	0.39	0.06	0.13	0.25	0.09
23	0.22	-0.02	0.13	-0.06	0.17	0.36	0.07	0.11	0.22	0.08
24	0.20	-0.02	0.13	-0.06	0.17	0.34	0.08	0.10	0.18	0.07
sum	4.68	-0.59	2.96	-2.61	3.11	8.25	1.33	6.44	3.65	1.63

Table 6.3

Appliance Load for Typical Hot Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.22	0.02	0.04	-0.05	0.15	0.30	0.02	-0.02	0.37	0.10
2	0.20	0.01	0.04	-0.07	0.15	0.26	0.03	0.0	0.32	0.09
3	0.17	0.0	0.03	-0.09	0.14	0.21	0.05	0.03	0.27	0.08
4	0.15	0.0	0.03	-0.11	0.14	0.18	0.05	0.06	0.23	0.06
5	0.14	-0.01	0.03	-0.13	0.13	0.16	0.06	0.11	0.20	0.06
6	0.14	-0.01	0.02	-0.14	0.12	0.16	0.05	0.15	0.19	0.06
7	0.15	-0.01	0.03	-0.14	0.10	0.19	0.04	0.19	0.19	0.06
8	0.16	0.0	0.03	-0.13	0.08	0.23	0.01	0.22	0.22	0.07
9	0.19	0.01	0.04	-0.11	0.06	0.29	-0.01	0.25	0.25	0.09
10	0.21	0.02	0.05	-0.09	0.04	0.35	-0.04	0.27	0.30	0.10
11	0.24	0.03	0.05	-0.08	0.04	0.40	-0.06	0.29	0.35	0.12
12	0.26	0.03	0.05	-0.07	0.04	0.45	-0.08	0.30	0.40	0.13
13	0.28	0.04	0.05	-0.07	0.05	0.47	-0.09	0.29	0.45	0.14
14	0.29	0.04	0.05	-0.08	0.07	0.48	-0.09	0.28	0.49	0.15
15	0.30	0.03	0.04	-0.09	0.09	0.48	-0.09	0.26	0.53	0.15
16	0.30	0.03	0.04	-0.10	0.12	0.47	-0.08	0.23	0.56	0.16
17	0.30	0.03	0.03	-0.11	0.14	0.45	-0.07	0.20	0.57	0.16
18	0.30	0.02	0.03	-0.11	0.16	0.44	-0.06	0.16	0.58	0.15
19	0.30	0.02	0.03	-0.10	0.17	0.43	-0.05	0.12	0.58	0.15
20	0.29	0.02	0.03	-0.09	0.18	0.42	-0.04	0.08	0.57	0.15
21	0.28	0.03	0.04	-0.07	0.17	0.40	-0.03	0.04	0.54	0.14
22	0.27	0.03	0.04	-0.05	0.17	0.39	-0.02	0.0	0.51	0.13
23	0.26	0.03	0.04	-0.04	0.16	0.37	-0.01	-0.02	0.47	0.13
24	0.24	0.02	0.04	-0.04	0.15	0.34	0.0	-0.03	0.42	0.11
sum	5.64	0.43	0.92	-2.16	2.82	8.31	-0.52	3.43	9.55	2.73

Table 6.4

95th Percentile of Appliance Load for Typical Cool Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.32	-0.12	0.28	-0.11	0.40	0.64	0.38	0.80	-0.12	0.05
2	0.28	-0.14	0.27	-0.14	0.39	0.50	0.41	0.83	-0.20	0.02
3	0.23	-0.15	0.25	-0.19	0.38	0.37	0.45	0.88	-0.28	0.0
4	0.20	-0.15	0.23	-0.25	0.37	0.26	0.49	0.96	-0.34	-0.02
5	0.17	-0.16	0.21	-0.31	0.35	0.20	0.51	1.05	-0.38	-0.03
6	0.16	-0.16	0.20	-0.35	0.33	0.21	0.50	1.15	-0.40	-0.03
7	0.17	-0.15	0.21	-0.37	0.32	0.31	0.47	1.24	-0.39	-0.01
8	0.20	-0.14	0.22	-0.37	0.30	0.48	0.42	1.33	-0.36	0.01
9	0.24	-0.12	0.24	-0.36	0.29	0.68	0.36	1.42	-0.32	0.04
10	0.28	-0.10	0.26	-0.34	0.28	0.88	0.30	1.49	-0.25	0.07
11	0.33	-0.09	0.27	-0.33	0.27	1.06	0.25	1.54	-0.18	0.10
12	0.36	-0.07	0.27	-0.33	0.27	1.19	0.22	1.57	-0.09	0.13
13	0.39	-0.05	0.27	-0.36	0.27	1.24	0.22	1.57	0.01	0.14
14	0.40	-0.04	0.25	-0.39	0.29	1.23	0.24	1.55	0.11	0.15
15	0.41	-0.03	0.23	-0.43	0.31	1.18	0.28	1.52	0.19	0.16
16	0.41	-0.03	0.22	-0.45	0.33	1.10	0.32	1.46	0.26	0.15
17	0.41	-0.03	0.21	-0.46	0.35	1.01	0.36	1.39	0.30	0.15
18	0.41	-0.03	0.20	-0.44	0.36	0.94	0.39	1.30	0.33	0.14
19	0.41	-0.03	0.21	-0.40	0.38	0.90	0.39	1.20	0.32	0.13
20	0.40	-0.05	0.23	-0.33	0.39	0.88	0.39	1.09	0.28	0.13
21	0.40	-0.06	0.25	-0.26	0.40	0.87	0.37	0.99	0.22	0.12
22	0.39	-0.07	0.27	-0.19	0.41	0.85	0.36	0.90	0.14	0.10
23	0.38	-0.09	0.28	-0.13	0.41	0.82	0.35	0.83	0.06	0.09
24	0.35	-0.11	0.29	-0.10	0.41	0.75	0.35	0.80	-0.03	0.07
sum	7.68	-2.15	5.79	-7.38	8.23	18.56	8.80	28.86	-1.11	1.87

Table 6.5

95th Percentile of Appliance Load for Typical Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.37	0.0	0.11	-0.09	0.37	0.62	0.25	0.51	0.39	0.11
2	0.33	-0.01	0.09	-0.12	0.36	0.48	0.29	0.53	0.31	0.08
3	0.29	-0.02	0.07	-0.17	0.34	0.35	0.32	0.58	0.25	0.06
4	0.26	-0.02	0.04	-0.23	0.33	0.24	0.35	0.65	0.20	0.05
5	0.24	-0.02	0.02	-0.28	0.32	0.18	0.36	0.74	0.17	0.04
6	0.23	-0.02	0.01	-0.32	0.30	0.19	0.36	0.83	0.16	0.04
7	0.24	-0.01	0.02	-0.33	0.28	0.29	0.33	0.93	0.16	0.05
8	0.27	0.0	0.03	-0.34	0.27	0.45	0.28	1.02	0.18	0.07
9	0.30	0.01	0.05	-0.33	0.25	0.65	0.22	1.11	0.22	0.10
10	0.34	0.03	0.07	-0.31	0.24	0.86	0.16	1.18	0.28	0.13
11	0.39	0.05	0.09	-0.31	0.24	1.04	0.12	1.24	0.35	0.16
12	0.42	0.06	0.10	-0.31	0.24	1.17	0.09	1.27	0.44	0.19
13	0.45	0.08	0.09	-0.33	0.24	1.22	0.08	1.27	0.54	0.21
14	0.47	0.10	0.07	-0.35	0.26	1.21	0.10	1.25	0.64	0.22
15	0.48	0.11	0.05	-0.38	0.27	1.15	0.13	1.20	0.74	0.22
16	0.48	0.12	0.02	-0.40	0.29	1.07	0.17	1.14	0.83	0.22
17	0.48	0.12	0.01	-0.41	0.31	0.99	0.20	1.06	0.89	0.22
18	0.48	0.12	0.0	-0.40	0.33	0.92	0.23	0.97	0.91	0.21
19	0.48	0.11	0.01	-0.35	0.35	0.88	0.24	0.87	0.90	0.20
20	0.47	0.10	0.03	-0.29	0.36	0.86	0.24	0.77	0.84	0.19
21	0.47	0.08	0.06	-0.23	0.37	0.85	0.23	0.68	0.77	0.18
22	0.45	0.06	0.09	-0.16	0.37	0.83	0.22	0.60	0.67	0.17
23	0.44	0.04	0.11	-0.11	0.37	0.80	0.22	0.54	0.57	0.15
24	0.41	0.02	0.11	-0.08	0.37	0.73	0.23	0.51	0.48	0.13
sum	9.23	1.13	1.37	-6.63	7.42	18.04	5.41	21.44	11.89	3.42

Table 6.6

95th Percentile of Appliance Load for Typical Hot Summer Day										
Hour	family	roomnum	refrig	apt	dish	dryer	stove	water	central	tv
1	0.42	0.10	-0.02	-0.08	0.33	0.60	0.16	0.29	0.77	0.15
2	0.37	0.08	-0.03	-0.12	0.33	0.46	0.20	0.32	0.68	0.13
3	0.33	0.07	-0.05	-0.17	0.31	0.33	0.24	0.38	0.61	0.11
4	0.30	0.07	-0.08	-0.23	0.30	0.22	0.27	0.45	0.55	0.09
5	0.27	0.06	-0.09	-0.28	0.28	0.16	0.29	0.54	0.52	0.08
6	0.26	0.07	-0.10	-0.32	0.27	0.17	0.28	0.63	0.50	0.08
7	0.27	0.07	-0.10	-0.34	0.25	0.27	0.25	0.73	0.51	0.09
8	0.30	0.08	-0.08	-0.34	0.23	0.43	0.20	0.82	0.53	0.12
9	0.34	0.10	-0.06	-0.33	0.22	0.63	0.14	0.91	0.58	0.15
10	0.38	0.12	-0.05	-0.31	0.21	0.84	0.08	0.98	0.64	0.18
11	0.43	0.13	-0.03	-0.30	0.20	1.02	0.03	1.03	0.72	0.21
12	0.46	0.15	-0.03	-0.30	0.21	1.15	-0.01	1.06	0.82	0.24
13	0.49	0.17	-0.04	-0.31	0.21	1.20	-0.01	1.05	0.93	0.25
14	0.51	0.19	-0.07	-0.34	0.23	1.19	0.0	1.02	1.04	0.27
15	0.52	0.21	-0.09	-0.36	0.24	1.13	0.02	0.97	1.15	0.27
16	0.53	0.22	-0.12	-0.38	0.26	1.05	0.06	0.90	1.25	0.27
17	0.53	0.23	-0.13	-0.38	0.28	0.97	0.09	0.81	1.31	0.27
18	0.53	0.23	-0.14	-0.37	0.30	0.89	0.11	0.72	1.34	0.26
19	0.53	0.22	-0.13	-0.33	0.32	0.85	0.12	0.62	1.32	0.25
20	0.52	0.20	-0.11	-0.27	0.33	0.84	0.12	0.53	1.27	0.24
21	0.52	0.18	-0.08	-0.20	0.34	0.83	0.12	0.44	1.19	0.23
22	0.50	0.16	-0.05	-0.14	0.34	0.81	0.12	0.37	1.08	0.22
23	0.48	0.14	-0.03	-0.09	0.34	0.77	0.12	0.31	0.97	0.20
24	0.45	0.11	-0.02	-0.07	0.34	0.71	0.13	0.29	0.86	0.18
sum	10.26	3.39	-1.75	-6.35	6.68	17.52	3.13	16.18	21.14	4.54

Figures 6.1 to 6.6 provide both the NWS and WS load for three sample characteristics: central air conditioning, water heating and family size. Each quantile and the mean estimates are provided for each characteristic. The NWS figures were drawn by estimating the impact of each characteristic during the "cool" week. The WS figures were estimated by subtracting the total effect during a "cool" week from the total effect during a "hot" week. Again, each of the quantiles and the mean estimates are shown. Note that the figures use different scales for kilowatt hours so care should be taken in their comparison.

Figure 6.1




Figure 6.2






Figure 6.3



Time (Sunday - Saturday)

Non-Weather Sensitive Water Heating Effect


Figure 6.4



Time (Sunday-Saturday)

Weather Sensitive Water Heating Effect


Figure 6.5



Time (Sunday-Saturday)

Non-Weather Sensitive Central Air Effect

Figure 6.6



Weather Sensitive Central Air Effect

Figures 6.1-6: Estimated Impact of Selected Characteristics on Conditional Quantiles (.25 to .95) and mean usage (M) for weather sensitive and non-weather sensitive components of the summer load measured in kilowatts. On the time axis 0 represents 0:00 a.m. on Sunday; 7 represents 24:00 on Saturday.

The tables and associated figures lead to several generalizations about the impact of individual characteristics on WS and NWS loads. Water heating and electric dryer have the largest impact on NWS load. Additional refrigerators also have a sizeable impact on the low use end of the load, but only a small impact on extreme values. Not surprisingly, the largest impact on weather sensitive load comes from central air conditioning. Moderate impacts occur as a result of additional room air conditioners, larger family sizes, more televisions and of living in a multi-family dwelling. Water heating and use of the electric stove are associated with decreases in the weather sensitive load.

While we did not have access to any prior studies that estimated end-use characteristics for the summer months, there are a number of studies that have attempted to estimate yearly or daily usage for various appliances by using data on *yearly* KWH usage. Eighteen of these studies are summarized in EPRI (1989), which includes a study of the Commonwealth Edison service area (Commonwealth Edison, 1985).

These estimates can be compared to the daily totals for a typical day in Table 6.1. Comparisons to our typical summer day give the following results. First, the values for refrigerator, stove, dishwasher, central air, water heater and number of television sets all fall within the range of values in the eighteen studies. Second, central air conditioning and room air conditioning are estimated to be below the Commonwealth value. However, the appropriate comparison for central air and for room air conditioners is probably somewhere between the typical and "hot" summer day. In this case our estimate for central air is higher than the Commonwealth estimate. The room air conditioning estimate, however, still remains low.* Third, our values for dryer and dishwasher are too high. The highest value for dryer in any of the studies is approximately 5 KWH/day, while our value is approximately 8 KWH/day. For some reason, our dryer results are apparently picking up the impact of omitted appliances. This phenomena is reported in the Commonwealth report to occur for their dishwasher variable (they arbitrarily divided their dishwasher estimate by four). Our dishwasher results are slightly above the highest estimate in four studies that included this variable. Finally, there were no comparable results for our characteristics of family size and apartment dwelling.

With the exception of refrigerators, extreme usage (e.g. 95% quantile estimates) is typically associated with an increase of approximately 200 to 300% in the estimates attributable to each appliance for most hours in the day. On typical or "hot" summer days, extreme use of electricity is primarily explained by increased use of central air, water heating or the dryer (probably standing in for other appliances). On "cool" summer days extreme use is again associated with water heating and use of the dryer, but is also associated with increases in use of the electric stove.

* Our estimate used number of room air conditioners while the other studies used existence of room air conditioning, which accounts for part of the difference.

7. Conclusions

We have suggested general methods for estimating nonparametric models of conditional quantile functions. What have we learned from using these methods to model the residential household demand cycle for electricity?

Our Stage I results indicate a very strong periodic component and weather impact on baseload (25% quantile) demand, even though this demand has relatively small variation over the cycle. This is primarily due to the fact that a large amount of data is used in estimation of these parameters, and the variation in these data is small. Thus, baseload demand can be accurately forecast for most customers and the variation throughout the summer in this baseload is small. On the other hand, there is much more variation in the extreme uses of electricity. This makes forecasts of this extreme use much more difficult for individual customers. Thus, while estimates of the 95% quantile have a strong periodic *shape*, variation around this shape is large.

In contrast, our Stage II results indicate that our ability to explain the variation in usage across customers typically increases as we move from baseload to extreme usage. There is little variation in baseload demand to explain across customers, and the covariates that we use are not well-suited to this task. There is, however, a great deal of variation in extreme use across customers and this variation is closely associated with many of the covariates that we use in this study.

In contrast to prior research in this area which focused exclusively on models for conditional mean parameters, our results strongly suggest that the statistical decomposition of electricity load curves into models for distinct quantiles can provide a rich source of additional detail.

Our results suggest that statistical decomposition of load curves can be especially useful for forecasting peak demands. Knowledge of changes in demographics can be of only minor help in forecasting summer baseload demand for individual customers. Thus the use of a

"representative customer" for estimating baseload residential demand can largely ignore changes in demographics and simply concentrate on the number of customers. However, estimates for capacity planning purposes that need peak load forecasts require close attention to changing demographics and the associated extreme usage of certain appliances and household characteristics. Finally, the difference between estimates of the conditional mean models and quantile models can be quite dramatic.

REFERENCES

- Antoch, J. and Janssen, P. (1989) Nonparametric regression m-quantiles *Statistics and Probability Letters* 8, 355-362.
- Commonwealth Edison Company, Strategic Analysis Department, Load Forecasting Section (1985) "Residential Unit Energy Consumption: A Conditional Energy Demand Approach", Xerox.
- Commonwealth Edison Company, (1986) "1985 RLS Appliance Survey", Internal Memo.
- Cox, D.D. (1983) Asymptotics for M-type smoothing splines, *Annals of Statistics*, 11, 530-551.
- DeBoor, C. (1978) *A Practical Guide to Splines*, New York: Springer-Verlag. Osborne, M. (1989) An effective method for computing regression quantiles, preprint.
- Engle, R. F., Granger, C., Rice, W. J., and Weiss, A. (1986) Semiparametric Estimates of the Relation Between Weather and Electricity Sales, 81, 310-320.
- EPRI, *Residential End-Use Energy Consumption: A Survey of Conditional Demand Estimates*, CU-6487, Research Project 2547-1, Palo Alto, CA: Electric Power Research Institute, October, 1989.
- Gutenbrunner, C., Jurečková, J. and Koenker R. (1990) A New Approach to Rank Tests for the Linear Model, preprint.
- Hall, P. and Sheather, S. (1988) On the distribution of a studentized quantile *Journal of the Royal Statistical Society (B)*
- Hastie, T. and Tibshirani, R. (1986) Generalized additive models, with discussion, *Statistical Science*, 1, 297-318.
- Hendricks, W., Koenker, R. and Podlasek, R. (1977) Consumption Patterns for Electricity, *Journal of Econometrics*, 5, 135-153.
- Hendricks, W., Koenker, R., and Poirier, D. (1979) Stochastic Parameter Models for Panel Data: An Application to the Connecticut Peak Load Pricing Experiment, *International Economic Review*, 20, 707-724.
- Jurečková, J. (1984) Regression Quantiles and Trimmed least squares estimator under a general design, *Kybernetika* 20, 345-357.
- Koenker R. (1987) A Comparison of Asymptotic Testing Methods for l_1 -Regression, *Statistical Data Analysis Based on the l_1 Norm*, ed. Yadolah Dodge, North-Holland.
- Koenker, R. and Bassett G. (1978) Regression Quantiles, *Econometrica*, 46, 33-50.
- Koenker, R. and Bassett G. (1982) Robust Tests for Heteroscedasticity Based on Regression Quantiles, *Econometrica*, 50, 43-60.
- Koenker, R. and d'Orey, V. (1987) Computing Regression Quantiles, *Applied Statistics*, forthcoming.

- Koenker, R. and Portnoy, S (1987) L-Estimation of Linear Models, *Journal of the American Statistical Association*, 82, 851-857.
- Koenker R. and Portnoy S. (1988) M-Estimation of Multivariate Regressions, Technical Report, University of Illinois.
- Lenth, R. (1977) Robust splines *Communications in Statistics: Theory and Methods* Bassett, G. and Koenker, R. (1986) Strong Consistency of Regression Quantiles and Related Empirical Processes, *Econometric Theory*, 2, 191-201.
- Lindley, D. V., and Smith, A.F.M. (1972) Bayes Estimates for the Linear Model, *Journal of the Royal Statistical Society (B)*, 34, 1-41.
- Mosteller, F. and Tukey, J.W. (1977) *Data Analysis and Regression* Reading, Mass: Addison-Wesley. Stone, C. (1977) Consistent Nonparametric Regression, with discussion, *Annals of Statistics*, 5, 595-645
- Poirier, D. (1973) Piecewise regression using cubic splines *Journal of the American Statistical Association*, 68, 515-524.
- Portnoy, S. (1990) Asymptotic behavior of regression quantiles in non-stationary, dependent cases preprint. Stone, C. (1985) Additive regression and other nonparametric models, *Annals of Statistics*, 13, 689-705.
- Portnoy, S. and Koenker, R (1989) Adaptive L-estimation of linear models, *Annals of Statistics*, 17, 362-381
- Ramsey, J.O. (1988) Monotone regression splines, with discussion, *Statistical Science* 3, 425-461
- Ruppert, D. and Carroll, R. (1980) Trimmed least squares estimation in the linear model, *Journal of the American Statistical Association*, 75, 828-838.
- Samanta, M. (1989) Non-parametric estimation of conditional quantiles, *Statistics and Probability Letters* 7, 407-412
- Sheather, S.J. and Maritz, J.S> (1983) An Estimate of the Asymptotic Standard Error of the Sample Median, *Australian Journal of Statistics*, 25, 109-122.
- Smith, A.F.M.(1973) A General Bayesian linear model, *Journal of the Royal Statistical Society Series B*, 35, 67-75.
- Troung, Y.K. (1989) Asymptotic properties of kernel estimators based on local medians, *Annals of Statistics*, 17, 606-617.
- Wahba, G., S. Wold(1975) Completely automatic french curve: fitting spline functions by cross validation, *Communications in Statistics*, 4, 125-141.
- Welsh, A.H. (1987) Kernel estimates of the sparsity function, in Y. Dodge (ed.) *Statistical Analysis Based on the L1 Norm*. New York: North-Holland.
- White H. (1990) Nonparametric estimation of conditional quantiles using neural networks, unpublished. Buja, A., Hastie, T. and Tibshirani, R. (1989) Linear smoothers and

additive models, with discussion, *Annals of Statistics*, 17, 453-555.

Appendix A

Consider a sequence of independent random variables $\{Y_i\}$ $i=1, \dots, n$, such that $Y_i \sim F_i$ and assume that θ^{th} conditional quantile of Y_i given a known vector x_i of covariates is linear in x_i , that is ,

$$F_i^{-1}(\theta) = x_i' \beta(\theta)$$

for some unknown parameter vector $\beta(\theta) \in \mathbb{R}^P$. The θ^{th} regression quantiles $\hat{\beta}_n(\theta)$ minimize

$$R_n(b) = \sum \rho_\theta(y_i - x_i b)$$

where $\rho_\theta(u) = (1 - \theta)|u|^- + \theta|u|^+$. One element of the potentially set-valued subgradient of $R_n(b)$ is

$$\nabla R_n(b) = - \sum \psi_\theta(y_i - x_i b) x_i$$

where $\psi_\theta(u) = \theta - \{u < 1\}$, and $\{A\}$ denotes the indicator function of the event A .

Write $y_i - x_i b = y_i - x_i(b - \beta(\theta)) - F_i^{-1}(\theta)$, set $\hat{\delta}_n = \sqrt{n} (\hat{\beta}_n(\theta) - \beta(\theta))$ and consider

$$g_n(\delta) = -\frac{1}{\sqrt{n}} \sum \psi_\theta(y_i - x_i \delta / \sqrt{n} - F_i^{-1}(\theta)) x_i$$

Clearly, $g_n(\hat{\delta}_n) \rightarrow 0$, by the definition of $\hat{\beta}_n(\theta)$ as a minimizer of $R_n(b)$. Using the argument of Ruppert and Carroll (1980) and others, for any $K > 0$

$$\sup_{\|\delta\| < K} \|g_n(\delta) - g_n(0) - E(g_n(\delta) - g_n(0))\| = o_p(1).$$

Thus, since

$$\begin{aligned} E g_n(\delta) &= -\frac{1}{\sqrt{n}} \sum [(\theta - 1)F_i(F_i^{-1}(\theta) + x_i \delta / \sqrt{n}) + \theta(1 - F_i(F_i^{-1}(\theta) + x_i \delta / \sqrt{n}))] x_i \\ &= -\frac{1}{n} \sum f_i(F_i^{-1}(\theta)) x_i x_i' \delta + o(1) \\ &\equiv -B_n \delta + o(1) \end{aligned}$$

we have the linear (Bahadur) representation

$$\sqrt{n} (\hat{\beta}_n(\theta) - \beta(\theta)) = \frac{1}{\sqrt{n}} B_n^{-1} \sum \psi_\theta(y_i - F_i^{-1}(\theta)) x_i + o_p(1)$$

and thus, $\hat{\delta}_n \xrightarrow{D} N(0, B_n^{-1} A_n B_n^{-1})$ where $A_n = \theta(1 - \theta)n^{-1} \sum x_i x_i'$.

NOTICE: Return or renew all Library Materials! The *Minimum Fee* for each Lost Book is \$50.00.


The person charging this material is responsible for its return to the library from which it was withdrawn on or before the **Latest Date** stamped below.

Theft, mutilation, and underlining of books are reasons for disciplinary action and may result in dismissal from the University.
To renew call Telephone Center, 333-8400

UNIVERSITY OF ILLINOIS LIBRARY AT URBANA-CHAMPAIGN

JUN 21 1991

1991

HECKMAN
NDERY INC. 
JUN 95
-To -Pleas® N. MANCHESTER,
INDIANA 46962

UNIVERSITY OF ILLINOIS-URBANA



3 0112 060295943